Data Replication in Low Bandwidth Military Environments – State of the Art Review

Dr. Allan Gibb*
Defence Research Establishment Valcartier
2459 Boulevard Pie-XI North
Val-Bélair, Québec, Canada G3J 1X5
418-844-4000, ext 4683
allan.gibb@drev.dnd.ca

Dr. Sam Chamberlain*
U.S. Army Research Laboratory
ATTN: AMSRL-IS-C
Aberdeen Proving Ground, MD, 21005-5067
410-278-8948
wildman@arl.mil

Abstract

Modern armies are undergoing a revolution in the way information is managed on the battlefield. Voice-based command, control, and communication systems are being complemented by, and in some cases replaced by (in whole or in part) digital command, control and communication systems. Digital systems offer the promise of increased battlefield awareness through a more systematic and automated distribution of relevant data than is possible with a voice-based communication system. To deliver on this promise, the communication backbone must be capable of distributing digital data among participating command and control nodes with no error and a timeliness appropriate to the operational scenario. To maintain information superiority, important information must be passed quickly enough to permit the friendly commander to stay within, and act within, the decision cycle of the enemy commander. On the tactical battlefield, the low data throughput and unreliable connectivity of the wireless communication links make it difficult to replicate enough data in a timely way to satisfy this objective. This paper reviews the state of the art of data replication mechanisms within a low bandwidth wireless military environment as revealed at a workshop sponsored by The Technical Cooperation Program (TTCP), Command, Control, Communications, and Intelligence (C3I) Group, Technical Panel 10 (TP-10) that was held at Fort Leavenworth, Kansas in April 1999.

1. Introduction

At the tactical level (below Brigade), units are highly mobile and have no choice but to use a broadcast medium (e.g., combat net radio) as the primary means of communication. The communication network is organized as a network of subnetworks, each subnetwork operating on a unique base channel. Subnetworks are linked through gateways that manage, among other things, the channel routing. Data can be broadcast from a given node to any other participating node on a subnetwork (single hop) or to one or more nodes on other subnetworks (multiple hops). It is also highly desirable, from the point of view of maintaining continuity of operations, to avoid a single point of failure in the information distribution system. A client-server model does not achieve this goal because, if the server fails for any reason, all client nodes dependent upon the server for information may become vulnerable. To avoid a single point of failure, the preferred approach is to replicate a shared situational picture among participating nodes (peer to peer replication). In the ideal world, each mobile node would maintain its own copy (in its local database) of its neighbor's situation information. This makes it possible for a node to assume the

* This study was conducted under the auspices of The Technical Cooperation Program (TTCP), Command, Control, and Communications (C3I) Group, Technical Panel 10 (TP-10) on Distributed Information Systems Technology.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **2000**	2. REPORT TYPE	3. DATES COVERED **00-00-2000 to 00-00-2000**
4. TITLE AND SUBTITLE **Data Replication in Low Bandwidth Military Environments - State of the Art Review**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Army Research Laboratory,ATTN: AMSRL-IS-C,Aberdeen Proving Ground,MD,21005-5067**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**		
13. SUPPLEMENTARY NOTES **2000 Command and Control Research and Technology Symposium**		
14. ABSTRACT		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**9**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

role of a neighbor at any time should that be necessary. It also permits a node that has been disconnected from the network for a period of time to seek missing information from any neighbor node upon reconnection. In the real world, with present-day radio technology, links are characterized by extremely limited data throughput (typically less than 2 kilobits/second; 300-600 bits/second is not uncommon) and unreliable connectivity. The low data throughput and unreliable connectivity of the wireless communication links on the tactical battlefield make it impossible to maintain perfectly synchronized copies of the information content in neighboring mobile databases. This means that the data in adjacent databases on the battlefield will not be fully consistent much of the time.

In the spring of 1999, a workshop was organized to address the problem of data replication among distributed databases using unreliable wireless communication links. The workshop was organized by The Technical Cooperation Program (TTCP)[1], C3I Group, Technical Panel 10 (Distributed Information Systems Technology) and was held at Fort Leavenworth, Kansas from 20-22 April 1999. The purpose of the workshop was to assess the state of the art of data replication and to exchange ideas as to how to use databases even when traditional consistency expectations are unrealistic. In other words, managing inconsistency may begin by accepting inconsistency. The workshop had 52 participants, and included military and scientific personnel from government laboratories, industry, and academe from the four participating TTCP countries (U.S., United Kingdom, Canada, and Australia). This paper reviews the state of the art of data replication mechanisms within a low bandwidth, wireless military environment as revealed at the workshop. See Appendix A for a list of the workshop presentations.[2]

2. Trends in commercial replication solutions

Replication mechanisms can be categorized from several perspectives, but the functional issues typically reduce to two basic characteristics: the resolution of the entities being replicated, and the frequency of replication, sometimes called consistency. Currently, two extremes exist in the resolution domain: *file replication* and *transaction replication*. Most commercial products address one of these two, but not both. Although the primary interest at the workshop was in transaction replication (or mechanisms that operate as part of a database management system), file replication is also a major concern for the military, especially within command posts. Replication frequency is also characterized by extreme requirements. At one extreme, every change must be propagated to maintain tight consistency between replicates. This expectation is common with transaction systems. At the other extreme, items are allowed to drift out of synchronization and are then periodically re-synchronized, usually rapidly using adequate bandwidth links. This approach is typical with file replication mechanisms. These relationships did not evolve for technical reasons, but are a result of the specific requirements of the application environments. Figure 1 illustrates these categories.

Figure 1 shows that file replication systems are currently geared towards an environment in which mobility is characterized by periodic resynchronization between file systems. This occurs as users check in and out of a network. Many changes to a file typically occur during the period of disconnection. In these systems an audit trail isn't maintain. How the files evolve to their

[1] For details of TTCP see http://www.dtic.mil/ttcp/.; for TTCP C3I group info see http://www.dtic.mil/ttcp/c3i.htm.
[2] To access presentations see http://www.dtic.mil/ttcp/.

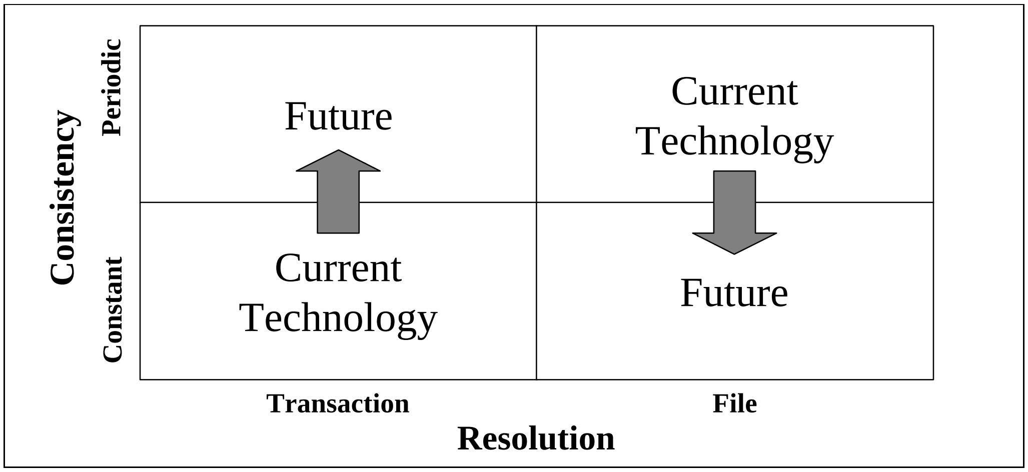

Figure 1: Categories of Replication Mechanisms

new state is unimportant —only the final state matters. Therefore, the update mechanism in file replication systems can be optimized to move efficiently from the initial to the final state without regard to intermediate states. In contrast, transaction replication systems (i.e., those within database management systems) have traditionally enforced strict consistency requirements. This entails that an audit trail be continuously and rigorously maintained so that failed transactions can be "rolled back" to a common, previous state. These systems are characterized by a constant and immediate attempt to maintain consistency between replicated entities. In the future, it is expected that these extreme characteristics will gradual merge into a single, continuous medium that will support variable replication schemes based on a variety of parameters. See [Chamberlain, 1996] for a discussion of promiscuous replication mechanisms.

2.1 *Enterprise Models – Commercial versus Military*

The enterprise model that is driving replication solutions for major database vendors still appears to be the client-server model. In this model, servers replicate data over high bandwidth links (typically Ethernet) using tight-consistency protocols (requiring a two-phase commit) while remote workers (clients) communicate with an enterprise server using wired or wireless links to upload/download information. In each case, disconnected operation is supported. However, the underlying assumption is that, upon reconnection, either the communications "pipe" is wide enough to support rapid resynchronization, or (in the case of low bandwidth wireless connections) latency is not an issue so the reconnected client can take as much time as required to resynchronize its database with the server.

An exception to this model is the Starburst[3] file replication system. It is optimized for an environment in which a single entity must be reliably sent to many destinations, such as a periodic update of a product inventory database. It uses a multicast file transfer protocol (MFTP)[4] for replication. MFTP accomplishes its optimization by unreliably broadcasting a complete file to all the recipients. However, the protocol encodes the information so that missed packets can be identified by each recipient. Each recipient then sends its list of missed packets back to the server. The server assembles a list of all missed packets and the complete list of all

[3] Recently acquired by Adero, Inc., see http://www.starburstsoftware.com/.
[4] Multicast File Transfer Protocol is now a trademark of Starburst Software, Inc.

missed packets is sent to all the destinations. This process is repeated until no hosts report any missed packets. The approach has been tested favorably over tactical radio nets in comparison to the standard, reliable, uni-cast FTP. This approach is optimized for a large number of users; clearly, the larger the destination set, the better the improvement over using many single uni-cast connections. However, the protocol does not address the problem of replication over tenuous communication with latency as a primary metric. That is a different domain of problems.

The aforementioned approaches do not adequately address the military, tactical battlefield problem where: (1) low bandwidth, wireless communications using a broadcast protocol is usually the only option, (2) continuous peer-to-peer replication is required to preserve operational autonomy of nodes, (3) disconnection is frequent and usually unplanned, and (4) latency is almost always an issue. There is a need for light-weight yet robust support for peer-to-peer replication (i.e., when each node provides its own replication service) designed to operate continuously over wireless broadcast communications media.

2.2 ***Wireless versus Mobile***

The terms “wireless” and “mobility’ are sometimes used interchangeably as if they are interchangeable terms - obviously, they are different. Often, it is mobility - the capability to disconnect and later re-connect and synchronize, that is the focus of the “wireless” domain described. Because this is a common occurrence in a wireless environment, mobile and wireless environmental characteristics are sometimes confounded. The disconnect/reconnect problem is well addressed in several research projects, but this is only a part of the wireless and/or low bandwidth problem. In the projects presented, it was assumed that the wireless (or other communications) environment could handle the load to resynchronize. On the tactical battlefield, this assumption is frequently invalid and is the source of significant frustration. One technique offered is *caching* where an the attempt is made to load what one needs before an expected disconnection occurs to minimize the need to download large data items at a later time. This is applicable to both file and transaction (database) replication approaches.

2.3 ***Adaptability***

Adaptability is the ability to actively modify information flow to changes in network performance and battlefield situations. It has received little attention from database management system vendors. To meet military needs, several capabilities must be included in replication mechanisms. First, network performance data must be available and no system adequately supports this requirement. Second, the data model must include a communications model so that there is a place to store and maintain network performance attributes. Third, there must be a way to make use of these data by the replication mechanism (e.g., improved trigger language semantics).

Three encouraging signs in commercial database evolution are worthy of note. First is emerging database vendor support for commercial, mobile users. Second is an increasing trend to support replication between heterogeneous databases. Third, as noted above, is the availability to users of measured transaction delays in output queues[5].

---

[5] A planned capability to do this was described by Sybase, Inc., see http://sybase.com/mec/papers.html for white papers.

3. Research Directions

The workshop provided encouraging signs that research is being carried out relevant to the topic of data replication in low bandwidth environments. A variety of approaches are being investigated. It was clear that the problem is multi-faceted and no single approach will provide all the answers.

3.1 *Procedurally-Based Information Management Techniques*

The essential purpose of managed data replication in a low bandwidth military environment is the intelligent selection of information to share. The low bandwidth constraint means that a constant, undisciplined, uncontrolled "push" of information is not supportable by the communication system. The constraint places a premium on careful management of information flow, both by application and network layers.

A key to information management in a low bandwidth environment is a set of procedural rules that define who on the battlefield should receive a given piece of information for a given situation. The military user (subject matter expert) must assist in establishing these rules that must, ultimately, be defined in a form suitable for automated application in a command and control system. The concept of triggers employed in active databases provides a promising mechanism for the implementation of such rules and should be exploited further. However, the long-term goal should be to rethink database concepts to move toward a model which unifies application logic and data storage into a single environment. In other words, rules for managing the data (including dissemination rules) should be seamlessly integrated with the data itself. In a low bandwidth context, the challenge is to achieve this goal while minimizing the quantity of management data that is transmitted with a data element when it is replicated. One aspect of this approach is to cease considering the database as a repository for fields of structured military messages (i.e., data elements) and to consider it as a repository for a dynamic model of the battlefield (model-based approach).

Another concept worth exploring is a *cost-based approach* to information dissemination. In this approach, the basis for controlling the information flow is a comparison between the communications "cost" of immediately sending the information versus the operational "cost" of *not* sending the information immediately, see [Wolfson *et al*., 1999]. Clearly, a key issue is the basis for assigning a numeric cost to each factor.

Better methods for simulating and measuring the impact of proposed information management techniques and data replication strategies are required so that new approaches can be evaluated before they are integrated into fielded systems. Latency alone will not reveal whether mission-critical information arrived on time or if serious information gaps are developing in a local situation picture. Tools are required for real-time monitoring of data consistency across replicated databases for selected domains. This unique military requirement is not well addressed by database vendors, network suppliers, or the simulation community.

3.2 *File and Transaction Replication*

For directory- and file-based replication systems, it will be necessary to minimize the amount of information exchanged when synchronizing directories on different machines. Latency

considerations must also be included in requirements for scalability and reliability. Presently complete files are transferred. When only changes in the files are transferred, then significant improvements can be obtained.[6]

For transaction replication systems, an important topic is database resynchronization upon reconnection. In this context, a major concern is conflict resolution. A conflict occurs when two users independently update the same data element in their local databases while disconnected, the values are not the same, and the replication process tries to reconcile the values when the databases are reconnected. Conflict resolution can demand considerable effort (and bandwidth) if user intervention is required. Automated conflict resolution schemes that minimize communication demands warrant investigation. The *version vector* approach described in Bengal[7] is an example of using time to resolve conflicts. A more general approach is needed for cases in which data quality and reliability are of equal concern. The role of different data ownership schemes in conflict avoidance is also an important research issue. In some cases, there are bandwidth and/or operational constraints that prevent a database from re-synchronizing completely upon reconnection. These topics are especially difficult to describe formally and require extending the tasks of managing database consistency to a particular user domain. By applying military science principles and expertise to the problem, consistency requirements may be relaxed to more realistic values, thus significantly reducing the number of resynchronization transactions required between different databases. This area of “gray synchronization” has not been adequately addressed.

A second important topic is minimizing uncertainty while disconnected through use of predictive algorithms. By exploiting information already in the database, such as last reported vehicle speed and direction, new values for other database elements may be derived via prediction (e.g., dead reckoning). Even better approximations can be achieved if planned route information is exchanged so that planned way points can be used for interpolation to further reduce errors and define realistic reporting thresholds.

Most of the dynamic relationships between data replicas presently supported by COTS relational database products are limited to table rows. This requires exchange of data fields within the rows that may not have changed. If the dynamic relationships could extend only to modified fields (i.e., selective replication) then a significant decrease in the quantity of information to be transferred would be obtained.[8]

To properly control information flow, the command and control nodes require continuous knowledge of network performance data, in particular, average session delay and average session throughput. It was exciting to hear that Sybase, Inc. is experimenting with measuring transaction delays in its output queues to make this information available to users.[9] It is recommended that more database vendors investigate this feature. To facilitate such a capability, the data model must include a communications model so that there is a place to store and maintain network

---

6 This planned capability was included in the presentation from Starburst Software, Inc.

7 Bengal is a product of Platinum Technology, Inc. that has since been acquired by Computer Associates, Inc. Version vectors are the basis of an approach called optimistic replication; see http://www.isi.edu/~johnh/PAPERS/Heidemann95d.html.

8 This technique is incorporated into the ATCCIS (Army Tactical Command & Control Information System) Replication Mechanism (ARM); see http://www.valcom-ottawa.com/software_development.htm.

9 This capability was described in the presentation by Tony Antonello from Sybase, Inc.

performance attributes. By storing this information in the database, it allows these parameters to be used by the replication mechanism and trigger language semantics.

4. **Impact of New Technologies**

Several of the presentations at the workshop dealt with replication of data among conventional databases residing in the application layer (command and control nodes). Conventional database technology is mature and well entrenched in the marketplace, in particular, relational technology. Several other technologies have the ability to influence, either directly or indirectly, how data replication will be carried out on the future tactical battlefield.

Middleware (CORBA, DCOM)[10] inserted between the application layer and conventional network layers, is based on distributed object technology and provides a set of services which can be very useful for transparently managing information flow in a distributed system. See [Anderson, *et. al*. 1997, 1999]. However, current versions of CORBA and DCOM require a level of communication overhead to support those services which is unacceptable in the low bandwidth regime (a few kilobits per second). The Object Management Group, or OMG[11], that champions CORBA is studying lightweight object request brokers (ORBs), but it does not appear to be a high priority topic and may wane before significant progress is accomplished.

Software agent technology may offer a new tool for managing database consistency. Intelligent agents installed on remote nodes may discover information of interest to a client node and request that this information be broadcast to the client node (and all other listening nodes). With this approach, only data known to be of interest to at least one node is replicated. This approach supports the intelligent selection of information to share and has the potential to reduce bandwidth requirements by limiting the quantity of data transmitted. Combining middleware and intelligent agent technologies may permit flexibility in the way artificial intelligence is applied to the problem because it removes the requirement that the intelligence be tightly integrated with the data store.

Communication technologies will also impact data replication on the future tactical battlefield. Clearly, technologies that can improve connectivity for distributed wireless nodes (satellite communications, cellular technology, airborne relays), improve data communications capacity (high capacity digital radios, asynchronous transfer mode (ATM)), or improve portable computer performance and capacity (personal digital assistants, palmtops, etc) will substantially facilitate the replication process. Also, the ability for protocols to share performance data between protocol layers, sometimes called "stack cognizance," can facilitate an ability to tune information and network management systems together, regardless of the available bandwidth; see [Todd, *et.al*., 1998].

5. **Conclusion**

This workshop was very successful in bringing together colleagues and collaborators in this rather specialized research domain. Government personnel (civilian and military), government contractors, database vendors and members of the academic community attended providing an

---

[10] CORBA: Common Object Request Broker Architecture; DCOM: Microsoft's Distributed Component Object Model; See [Orfali *et al*., 1996] for a tutorial on these technologies.

[11] See http://www.omg.org/ for information on OMG programs and activities.

eclectic mix of government, industry and academic perspectives. The database vendors presented their approach to data replication. In return, they were sensitized to the military-specific aspects of data replication that were not well supported by their products. All parties left with a heightened awareness of the research challenges posed by the need to replicate data in a low bandwidth wireless military environment.

Although the workshop did not greatly alter the direction of military research, it did provide optimism that industry is beginning to recognize the problems encountered in the military environment. It appears that new emphasis is being placed on mobile, wireless environments and that techniques are being developed and incorporated into commercial products to help alleviate the difficulties associated with imperfect communications.

6. References

[Anderson *et al*., 1997] Les Anderson, Sam Chamberlain, Richard Metzger, Tony Newton, John Laws, John Mansfield, Tony Gale, John Robinson. *Trends and Applications in Common Operating Environments and Distributed Computing Environments.* Report from the 1996/1997 TTCP STP-10 Study And Workshop, Published by The Communications Research Centre (CRC) and The Defence Research Establishment, Ottawa (DREO), Ottawa, Canada, July 1997.

[Anderson *et al*., 1999] Les Anderson, Sam Chamberlain, Eric Dorion, Carl DeFranco, John Laws, Iain Macleod, Tony Rathmell, John Robinson. *The Requirements for Middleware to Support C3 on the Move*, TTCP C3I Group TP-10 Workshop Report published as Defence Evaluation and Research Establishment (DERA) Report, July 1999.

[Chamberlain, 1996] Sam Chamberlain. "Resilient Data Replication Mechanisms for Battle *Proceedings of the 1996 Command and Control Research and Technology Symposium;* Naval Postgraduate School; 25-28 Jun 1996; pp 436-446. See http://www1.arl.mil/~wildman/PAPERS/papers.html.

[Orfali *et al*., 1996] Robert Orfali, Dan Harkey, Jeri Edwards. *The Essential Distributed Objects Survival Guide*. John Wiley & Sons, Inc., New York, 1996.

[Todd et al., 1998] Charlene Todd, Adarshpal Sethi, and Sam Chamberlain. "Feedback of Network Performance To Provide Better Utilization of Communication Resources," *Proceedings of the 2nd Annual Fedlab Symposium*, College Park, MD, 2-6 February 1998. See http://atirp.isr.umd.edu/.

[Wolfson *et al*., 1999] O. Wolfson, P. Sistla, B. Xu, J. Zhou, S. Chamberlain, N. Rishe, Y. Yesha. "Tracking Moving Objects Using Database Technology in DOMINO,", *Proceedings of NGITS'99, The Fourth Workshop on Next Generation Information Technologies and Systems* (Springer-Verlag Lecture Notes in Computer Science, Number 1649), Zikhron-Yaakov, Israel, July 1999, pp. 112-119.

APPENDIX A: List of Presentations

Title	**Speaker**	**Company / Organization**

Tuesday, 20 April 1999

Title	Speaker	Company / Organization
Resilient Replication Mechanisms	Dr. Sam Chamberlain	US Army Research Lab.
The CODA File System	Prof. Mahadev Satyanarayanan	Carnegie-Mellon University.
BENGAL	Mr. Todd Ekenstam	Platinum Technology, Inc.
Tactical Information Exchange Requirements	Major R.K. Ferguson	Canadian Forces LO, Ft. Gordon
Tracking Moving Objects	Prof. Ouri Wolfson	University of Illinois - Chicago
Testbed for the Evaluation of Battlefield Information Management Techniques Applied to a Low Bandwidth Tactical Wireless Communications Environment	Dr. Allan. Gibb & Jean-Claude St-Jacques	Defence Research Establishment Valcartier, CA
Informix Enterprise Replication	Mr. Madison Pruet	Informix, Inc.
Experiences from MCS	Mr. Joe Gilchrist	Lockheed-Martin
Data Consistency Issues in Support for Deployed Headquarters	Dr. Iain Macleod	Defence Science & Tech Org. C3 Research Centre
The TRIERARCH Trigger Architecture - Practice and Experience with Java and Other COTS Technologies	Mr. Richard Rabbat	Mass. Inst. Of Technology

Wednesday, 21 April 1999

Title	Speaker	Company / Organization
Next Generation Command and Control: Introducing the Artillery Regimental Data System (ARDS/ADM)	Mr. Michael Jones	MacDonald-Dettweiler
Replication and Reconciliation in Mobile Wireless Databases	Prof. B. R. Badrinath	Rutgers University.
Starburst Omnicast Concepts and Multicast FTP	Mr. Richard O'Brien	Starburst, Inc.
Data Streamlining	Major R.K. Ferguson	Canadian Forces LO, Ft. Gordon
Model-Based Battle Command and Change Propagation in Advanced Database Systems	Mr. Gautam Thacker	Lockheed-Martin Advanced Technology Lab
Technical Concepts for the ATCCIS Data Replication Mechanism	Mr. Tony Antonello	Sybase, Inc.
Sybase Strategies for the Replication and Synchronization of Data	Mr. Tony Antonello	Sybase, Inc.
Middleware Issues for Low Bandwidth, Mobile Wireless Environments	Dr. Ravi Jain	Telcordia Technology
Situation Awareness System for Canada	Mr. Andrew Hill	Computing Devices, CA
The Joint Common Database Replication System	Mr. Bob Carneval	US Army Program Executive Office for C3 Systems